

LSI 01-022

PATENT

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR PATENT

ON

*METHOD FOR CHANGING FIBRE CHANNEL SPEED OF A DRIVE LOOP WITH
ESM-CONTROLLED DRIVE BOXES USING REDUNDANT DRIVE CHANNELS*

BY

JAMES A. LYNN
1003 TIFFANY COURT
ROSE HILLS, KS 67133
CITIZEN OF USA

JEREMY D. STOVER
2803 N. SHEFFORD CIR.
WICHITA, KS 67205
CITIZEN OF USA

JASON M. STUHSATZ
2327 VENTOSO
WICHITA, KS 67204
CITIZEN OF USA

DAN A. RIEDL
425 W. DOUGLAS
ANDOVER, KS 67002
CITIZEN OF USA

TIMOTHY FLYNN
1955 N. MARKET
WICHITA, KS 67214
CITIZEN OF USA

CERTIFICATE OF MAILING BY "EXPRESS MAIL"

"Express Mail" Mailing Label Number: EV 013 245 333 US

Date of Deposit: October 31, 2001

I hereby certify that this correspondence is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 C.F.R. § 1.10 on the date indicated above and is addressed to Box Patent Application, Assistant Commissioner for Patents, Washington, D.C. 20231

BY: Penny L. Flint
Penny L. Flint

*METHOD FOR CHANGING FIBRE CHANNEL SPEED OF A DRIVE LOOP WITH
ESM-CONTROLLED DRIVE BOXES USING REDUNDANT DRIVE CHANNELS*

FIELD OF THE INVENTION

5 [0001] The present invention generally relates to the field of data communications storage, and particularly to a bridge controller used in conjunction with dual-ported fibre channel disk drives that have multi-speed capability and are physically organized into multiple enclosures, each of which is managed by an enclosure services component.

BACKGROUND OF THE INVENTION

10 [0002] The fibre channel is a serial data communication network architecture and protocol for interconnecting computers and peripheral devices. The fibre channel supports several protocols, including the Small Computer System Interface (SCSI) protocol, an ANSI standard for controlling peripheral devices by one or more host
15 computers. The computer or peripheral devices are connected to the network through fibre channel ports or other means. The host contains the fibre channel port. The fibre channel port uses a local data bus, such as a peripheral computer interface (PCI) bus, to implement data transfers.

[0003] Because the high bandwidth and connection flexibility it offers, fibre channel is
20 often used for connecting peripheral devices within multi-disk drive enclosures, such as redundant arrays of inexpensive disks (RAIDs), and for connecting multi-disk drive enclosures with one or more host computers. These multi-disk drive enclosures offer great economy, greater storage, and redundancy that improve operational reliability.

[0004] Malfunctioning enclosures can, in certain cases, degrade or disable
25 communications. There is the possibility of inserting a disk drive that is incapable of operating at the speed its enclosure has established for its remaining drives. There are no physical limitations which would prevent this event from occurring.

[0005] In the prior art, it is up to the enclosure services component to detect that the drive is not communicating properly and effectively remove it from the fibre loop by asserting

its bypass control line. This has the effect of presenting information that indicates that the drive has failed without any indication that the reason the drive has failed is only because of its mismatched speed.

[0006] Therefore, it would be desirable to provide a method for changing disk drive speeds in a controlled manner. It would also be desirable to provide such a method with simplified circuitry.

SUMMARY OF THE INVENTION

[0007] Accordingly, the present invention is directed to an apparatus having disk drives which are physically organized into multiple enclosures and a method for adjusting the speed of the disk drives. In a first aspect of the present invention, a method is employed for changing the speed of a channel on the enclosure services channel.

[0008] The present invention discloses a method for adjusting the speed of operation of a channel for communicating with disk drives in a multi ported system organized into a plurality of enclosure services modules and having a first channel connected in sequence from a bridge controller to a first enclosure services module and successively connected to successive enclosure services modules to a last enclosure services module and a second channel connected in reverse sequence from the bridge controller to the last enclosure services module and successively connected to the successive enclosure services modules to the first enclosure services module, the steps comprising: changing the speed of either the first or second channel on the bridge controller and using the other of the first and second channel to send the change speed frame from the bridge controller to each of the plurality of enclosure services module.

[0009] The present invention discloses a system for adjusting the speed of operation of a channel for communicating with disk drives in a multi ported system, comprising a bridge controller having a first channel and a second channel and a plurality of enclosure services modules, each having a first channel connected in sequence from a bridge controller to a first enclosure services module and successively connected to successive enclosure services modules to a last enclosure services module and each having a second

channel connected in reverse sequence from the bridge controller to the last enclosure services module and successively connected to the successive enclosure services modules to the first enclosure services module

[0010] It is to be understood that both the forgoing general description and the following detailed description are exemplary and explanatory only and are not restrictive of the invention as claimed. The accompanying drawings, which are incorporated in and constitute a part of the specification, illustrate an embodiment of the invention and together with the general description, serve to explain the principles of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

[0011] The numerous advantages of the present invention may be better understood by those skilled in the art by reference to the accompanying figures in which:

FIG. 1 illustrates a system diagram.

FIG. 2 illustrates a flowchart of the steps taken in changing the disk speed of the enclosures.

DETAILED DESCRIPTION OF THE INVENTION

[0012] The present invention relates to a multi-speed algorithm which uses the bridge controller as the central point of speed control.

[0013] Reference will now be made in detail to the presently preferred embodiments of the invention, examples of which are illustrated in the accompanying drawings.

[0014] Referring generally now to FIGS. 1 through 2, exemplary embodiments of the present invention are shown.

[0015] The fibre channel is a serial data communication network architecture and protocol for computers, workstations, and peripheral devices. The fibre channel port includes serial transmitter and receiver components coupled to a communications medium via a link that has wires made of a signal transmitting material. It could be metal wire, optical fiber, or some other transmissive material.

[0016] In fibre channel, data is transferred one bit at a time at extremely high transfer rates. The smallest unit of data, or grouping of data bits, supported by an fibre channel network is a 10-bit character. Fibre channel primitives are made up of 10-byte characters or bytes. Certain fibre channel primitives carry control information exchanged between fibre channel ports. The next higher level of data organization, fundamental to fibre channel protocol, is the frame. A frame may be formed of between several to several thousand bytes of data according to the type of data included. The fibre channel protocol specifies a next higher organizational level called the sequence.

[0017] A physical layout of an enclosure partitioned disk drive system is shown in FIG. 1. The enclosure services module and its drives constitute the enclosure. In a data processing system having a data processor coupled to a channel, such as an SCSI channel, disposed for transmitting and receiving data between the data processor and a peripheral storage subsystem, a bridge controller for expanding the maximum allowable number of disk drives is connectable to the channel. The bridge controller comprises a microprocessor having input/output terminals coupled to a CPU bus; a memory; a buffer; and additional circuitry coupled to the disk drives; and, a circuit for controlling the direction of data flow through the buffer. The circuit has input/output terminals coupled to the CPU bus and a control output coupled to a direction control input terminal of the buffer. The bridge controller is a target controller that uses the channel to connect to the peripheral device. Bridge controllers are commonly used in RAID arrays.

[0018] An in port and an out port of channel ch 1 of the bridge controller is connected by cable to the respective first pair of input and output ports 71 of a first enclosure services module 20 through which signals of channel pass through to the corresponding second pair of input and output ports 81 of the first enclosure services module 20. The second pair of input and output ports of the first enclosure, in turn, are connected to a pair of input and output ports of a second enclosure. The connections are made in such a manner until all the enclosure services modules are connected to channel 1.

[0019] The electrical connections are made, in reverse order, for channel ch 2. In this case, the bridge controller's channel 2 is connected first to the last module which was

connected for channel ch 1. Then, the next to the last module to be connected is connected to the last module. As shown in FIG. 1, the first pair of ports 70 of channel ch 2 are internally connected to a bypass 60. Bypass 60 can route the signals through a set of bypasses 50 which feed the signals to the disk devices 30. SES 40 is an SCSI device that monitors and controls enclosure services to enable the enclosure services module to perform certain functions such as channel speed changes and bypassing logic.

[0020] The present invention relates to a multi-speed algorithm which uses the bridge controller as the central point of speed control. To change speed, the bridge controller performs the following steps.

10 [0021] Various frame formats known from the prior art may be used. The enclosure services module is responsible for changing the speeds of all of its drives when it receives the change speed frame. Driving the speed control pins to the new speed value does this. The drive is responsible for recognizing, on the fly, that the speed has changed.

15 [0022] This algorithm works if there are at least two channels from the bridge controller to the enclosure services modules. The wiring is as depicted in FIG. 1. One channel is wired in reverse of the other channel(s). When an enclosure services module receives a change speed frame, it informs the other enclosure services modules on the same enclosure, so the speed setting of the entire enclosure is changed.

20 [0023] The ESM determines the speed capacity of a disk drive by trying two or more speeds and then checking to determine if the drive operates properly at the various speeds.

[0024] FIG. 2 illustrates the process steps in making a disk drive speed change. The bridge controller 10 changes the speed on the bridge controller 110 on channel ch 2. A change speed frame is sent to the last enclosure services module on channel ch 1 120. The last enclosure services module receives the speed change frame and changes its internal speed. The bridge controller waits for this enclosure to come up on channel ch 2. The next to the last enclosure services module is sent a speed change frame by bridge controller 10. The next to the last enclosure services module receives the change speed frame 130 and changes its internal speed. The bridge controller 10 waits for the

enclosure to come up on channel ch 2. The bridge controller determines whether this was the first enclosure services module 140. If it was, then the bridge controller changes the speed on the bridge controller 160 on channel 1, and waits for all of the enclosures to come up on that channel, and processing stops 170. Otherwise, the second to the last enclosure services module is selected 150. This second to the last enclosure services module is sent a change speed frame 130 and changes its internal speed. The bridge controller waits for the enclosure to come up on channel ch 2. The processing continues until the first enclosure is reached and processed.

[0025] The advantage of this approach is that the change is done in a controlled manner. One of the problems inherent in integrating a fibre channel solution is that the control of some behavior is distributed across the components. This, of course, can work fine, and, in many cases, is necessary. However, it does create the potential for interoperability problems between components. When these problems arise, the state of the system can become unstable.

[0026] As the bridge controller provides centralized control, the invention does not suffer interoperability problems.

[0027] It is believed that the system and method for changing fibre channel speed of a drive loop with ESM-controlled drive boxes using redundant drive channels of the present invention and many of its attendant advantages will be understood by the forgoing description. It is also believed that it will be apparent that various changes may be made in the form, construction and arrangement of the components thereof without departing from the scope and spirit of the invention or without sacrificing all of its material advantages. The form herein before described being merely an explanatory embodiment thereof. It is the intention of the following claims to encompass and include such changes.